20th EURO Working Group on Transportation Meeting, EWGT 2017, 4-6 September 2017, Budapest, Hungary

Inferring the actual urban road environment from traffic sign data using a minimum description length approach

Zoltán Fazekas*, Gábor Balázs, László Gerencsér and Péter Gáspár

Institute for Computer Science and Control (MTA SZTAKI),
Budapest, H-1111, Kende u. 13-17, Hungary

Abstract

In our paper, we focus on a group of traffic signs and use traffic sign logs to statistically infer the type of urban environment in which a car is being driven. The traffic signs are either perceived and logged by a human data entry assistant, or preferably automatically detected and logged by an on-board traffic sign recognition and logging system. An entry in the log-file records the traffic sign type and the along-the-route location of the sign. Furthermore, in case of collecting training data it records also the actual road environment category entered by the data entry assistant. The logs are seen as realizations of an inhomogeneous marked Poisson process, and the minimum description length (MDL) principle is applied to infer the actual environment. The aim of this approach is to encode the current data in the shortest possible way – assuming stochastic models derived from data collected earlier – and thereby accept the corresponding model and environment as actual. To evaluate the quality of classification, the inferred environment categories are compared to the ground truth data.

© 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the scientific committee of the 20th EURO Working Group on Transportation Meeting.

Keywords: Traffic sign recognition systems; Detection of road environment; Minimal description length principle; Marked Poisson point process; Statistical inference.

* Corresponding author. Tel.: Tel.: +36 1 2796163; fax: +36 1 4667483.
  E-mail address: zoltan.fazekas@sztaki.mta.hu

1. Introduction

Car drivers are assisted in many ways in perceiving and understanding the conditions of driving. This assistance is provided to them via road and traffic signs, information and directional signs, printed and electronic maps and navigational devices, on-board sensors, and community-based information services. For the evaluation of the importance of the various means of driver assistance, see Piao and McDonald (2008) and Amditis et al. (2010). The drivers of smart cars are further helped by their cars' advanced driving assistance systems (ADAS). Such systems comprise various high-end in-vehicle systems, e.g. anti-lock braking, lane-departure warning and traffic sign recognition systems (TSR) as subsystems. These subsystems are designed – separately and as a complex system – to increase the safety, and particularly the driver and the passenger safety, in the traffic, as well as the driver comfort. A recent ADAS architecture that implements numerous smart functions and comprises characteristic ADAS subsystems is presented by Fernandes et al. (2014).

Herein, traffic sign data – possibly gathered by some TSR system – is examined in conjunction with a particular aspect of driver assistance. More concretely, a statistical approach is presented which infers – from the log of traffic signs and from their distribution along a route – the type of the urban environment the car is being driven in, and detects the transitions between these. Before going into detail in presenting the approach, some motivation should be given concerning the following questions:

- Why is it important to categorize the road environment, and particularly, why is it important to categorize the urban road environment?
- Why should one use traffic sign logs to do so?

In different road environments, there are different things to look out for and to be aware of. For example, one should drive more cautiously in a densely populated busy downtown area than, say, in a calm residential area with virtually no traffic. Also, one should be more prepared to see and tolerate a relatively slow and steady flow of heavy vehicles in industrial and business areas (e.g., along long stretches of road near factories, or near large supermarkets) than in a downtown environment.

In conjunction with the risks associated with driving in various environments, a pilot study was carried out by Fazekas et al. (2012) concerning the abrupt braking and steering events – associated with trucks – in different socio-cultural environments. The results indicate that different driving styles – manifesting themselves in different average speeds before braking and different average intended decelerations – are exercised in different socio-cultural environments by the drivers. One could, however, argue that these differences in driving are partly due to certain traffic signs installed within these environments. Still, it would be probably worthwhile for the road vehicle manufacturers to provide assistance to the drivers also in this respect in the form of a dedicated ADAS function.

Why should – or at least could – one use logs of traffic signs to find out the type of the road environment they currently drive their car in? Before answering this question let us quote the following evaluation of the development efforts that appears among the concluding remarks of the paper – written by Møgelmose et al. (2012) – surveying vision-based TSR systems for assisted driving.

"Many contributions cite driver assistance systems as their main motivation for creating the system [i.e., a TSR system], but so far, only little effort has gone into the area of combining TSR systems with other aspects of driver assistance, and notably, none of the studies include knowledge about the driver's behaviour to tailor the performance of the TSR system to the driver."

Though not quite this kind of effort, but a very similar one is presented herein. Instead of reinforcing the ADAS subsystems and increasing the reliability of their outputs via fusing information from other ADAS subsystems, the re-use of the TSR computations and of the TSR output – in an aggregated manner – for road environment detection (RoED) is proposed. RoED can facilitate the computer vision and image understanding computations – carried out in a TSR system – by rendering the traffic sign detection more focused (e.g., by providing geometrical constraints and regions of interests) and, as a consequence, make the traffic sign recognition more robust and reliable. On the other hand, a TSR-based RoED system could be seen as a low-cost surrogate for a comprehensive vision-based road environment understanding and recognition system. A relatively low-cost RoED solution can be reached in this manner as many smart production cars – particularly in and above the medium price range – feature built-in TSR systems. Such TSR systems could be easily modified by the manufacturers to log the detected traffic signs and use this log to infer the type of the actual urban environment.

2. Traffic signs and road environments

The research effort communicated herein is the continuation of the work by Fazekas et al. (2016). There, a conceptually similar approach for detecting the transition between different topographical road environments – when driving from one environment to the other – was presented. The change detection task addressed there was admittedly not very practical. It was chosen as an easy-to-verify test case for the approach (i.e., for using traffic sign logs as input for change detection in the road environment) and methodology (i.e., the application of the minimum description length principle). In the present paper, the change of the urban scenery – often experienced when driving from the periphery of a town to its centre, or the other way round – is traced from the data recorded in traffic sign logs. Here, again the intention is to identify the change-point between environments; in this case, between certain urban road environments. Though the approach taken and the methodology utilized are common in the two communications, the aim is clearly more meaningful in the present case, furthermore, the task at hand is somewhat more complicated.

Relying on statistical inference in the given context is motivated by the fact that certain traffic signs appear more frequently in downtown than elsewhere; these signs include the ones indicating railway/bus station, restaurant, hotel/motel, cafeteria/refreshments, museum/historic building, parking places, particularly parking places against fee in the vicinity, and the ones warning the drivers of pedestrian traffic. While other traffic signs appear more frequently in suburbs (e.g., traffic signs indicating industrial area, goods harbour, airport, low-flying aircraft/sudden aircraft noise), or in rural areas (e.g., traffic signs warning of cattle/wild animals, and falling rock). Though, the afore-mentioned traffic signs associated with downtown areas appear mostly there, they are not necessarily the best choice for identifying an urban area as downtown. This is because one needs to come across a railway/bus station, a restaurant, or a museum – or at least the corresponding traffic signs – before deciding the area category. More frequent – but seemingly not that characteristic – traffic signs, such as the ones shown in Fig. 4, are therefore seen as better choices for the purpose.

3. Mathematical background

3.1. Marked point processes

A convenient model for describing traffic sign data is a marked point process. A point process is customarily given by an increasing sequence of time points, say T_n. However, the traffic signs that appear one after the other along a route are perhaps better described in space than in time. For this reason, the driving distance, or path-length, rather than time was chosen to characterize the point process corresponding to traffic sign data. The points of a point process may be labeled with marks. A marked point process then can be formalized as a pair (T_n, ρ_n), where ρ_n is the mark. For instance, in a log of traffic signs, a traffic sign location – expressed with the path-length of the route – may carry a label stating the type of the sign. In many practical cases that involve marked point processes, the marked Poisson processes have proved convenient and flexible. Although, the Poisson process is a continuous-time (or continuous space) model, in the algorithm presented here its discrete-time, or more precisely discrete space, approximation was applied.

3.2. Change detection

The problem of detecting abrupt changes in the dynamics of stochastic signals has been widely discussed in the literature together with important applications. Initially, the change detection within independent and identically distributed random data was the main target of research. This effort led to the well-known Page-Hinkley change detector (PHCD), see Page (1954), Hinkley (1971), Lorden (1971). The PHCD was later adopted and analyzed also for dependent data. See references in this regard and concerning other developments in Fazekas et al. (2016). The most important performance criteria for a change detector are its average run length between false alarms – which translates to false alarm rate – and the expected delay in detection.

3.3. Minimum description length approach in model selection and change detection

A novel approach to change detection based on the minimum description length (MDL) principle – the MDL was proposed in Rissanen (1978), and later extensively developed in Rissanen (1998) – was suggested in Baikovicius and Gerencsér (1990) and elaborated in Baikovicius and Gerencsér (1992). The MDL principle has its theoretical foundations in information theory. The stochastic approach based on this principle is used with success in tasks, such as model selection, feature extraction, as well as certain summarizing tasks, see Fua and Hanson (1991), Gao et al. (2000), Lakshmanan et al. (2002), and Kiernan and Terzi (2009) for examples and further references. The basic idea of the MDL approach is to choose between models for describing data on the basis of the minimum code-length by which one can encode the data relying on these models. The advantage of the MDL methodology is its enormous flexibility; for instance, the widely used PHCD can be interpreted as a procedure relying on this approach.

3.4. The Page-Hinkley change detector

Assume that we have a sequence of observations $\xi_1, \dots, \xi_N$, which is composed of two parts. The first part of the sequence is an independent identically distributed (iid) sequence of random variables taking discrete values according to the probability law $p(\xi_n, \theta_1)$, while the rest of the sequence is generated according to another probability law $p(\xi_n, \theta_2)$. The problem is then to estimate the time – or the location – of the change between the two probability laws from observed data in a real time manner. An MDL approach to solve this problem is as follows. Choose an arbitrary time – or location – τ, and assuming that this is when – or where – the transition between the probability laws takes place, encode the observed data optimally using hypotheses concerning the data generating mechanism. Based on the standard results of information theory, the overall optimal code-length $L_N(\tau)$ of the observed data in an asymptotic sense and allowing block coding is

$$L_N(\tau) = \sum_{n=1}^{\tau-1} -\log p(\xi_n, \theta_1) + \sum_{n=\tau}^{N} -\log p(\xi_n, \theta_2) \quad . \tag{1}$$

Following the MDL principle, the estimator of the transition time/location is obtained by minimizing the overall (assumedly) optimal code-length in τ. A heuristic procedure for minimizing L_N in real-time is obtained by identifying the time-point – or the location – after which L_N has a definite upward trend. Reformulating this characteristics of L_N, see Fazekas et al. (2016), one gets a sequence which is typically 0 before the true change-point, and typically increasing after that. This is actually the output of the well-known PHCD – see references in Subsection 3.2. – when the detector is applied to detecting the change in the parameter of the probability law.

We just quote here the formula derived in the paper cited above for the discrete time/discrete location approximation of a marked Poisson process. It was used for producing the results and the diagrams in Subsection 5.1. It was obtained by assuming an inhomogeneous iid sequence of random variables with binomial distribution, taking values 1 and 0, with probabilities θ_i and $1 - \theta_i$, respectively. Here, index i is the sequential number of the model (or in the given case, that of the road environment, e.g., downtown area: 1, residential area: 2 and industrial/commercial area: 3). The difference of the optimal code-lengths – often called score – encoding the j-th observation using the two respective probability laws is

$$\Delta L(j) = -\xi_j \cdot \log \frac{\theta_1}{\theta_2} - (1 - \xi_j) \cdot \log \frac{1 - \theta_1}{1 - \theta_2} - \sum_{k=1}^{m} \xi_j \cdot \zeta_{j,k} \cdot \log \frac{p_{1,k}}{p_{2,k}} \quad . \tag{2}$$

Here, probabilities θ_1 and θ_2 pertain specifically to downtown areas and residential areas, respectively. Similar formulas can be derived for other pairs of road environments. Relying on these formulae, one can compute the overall optimal code-length (assuming a given pair of road environments); it is done by applying the PHCD to find the estimator of the change-point τ. The specific signal to be monitored by the detector is derived from $\Delta L(j)$'s defined in (2). It is properly developed in the cited paper. The outputs of PHCDs used for detecting change in the road environment based on the collected traffic sign data are shown in Figs. 5 - 7 as examples.

4. Data collection from urban environments

For the purpose of the pilot study presented herein, a car-based data collection was carried out in respect of traffic signs and urban environments in three urban areas in Hungary. The areas involved in this study were

- Csepel — a district of Budapest situated on the Csepel Island of the river Danube, which had been – until about the end of 1980's – an industrial area with big factories producing mostly heavy industry products,
- Vác — a town on the east bank of the river Danube featuring a historic downtown and peripheral areas with either industrial/commercial, or residential buildings, and
- Százhalombatta — a town on the west bank of the river Danube featuring an electrical power station and an oil refinery, as well as housing projects and residential areas.

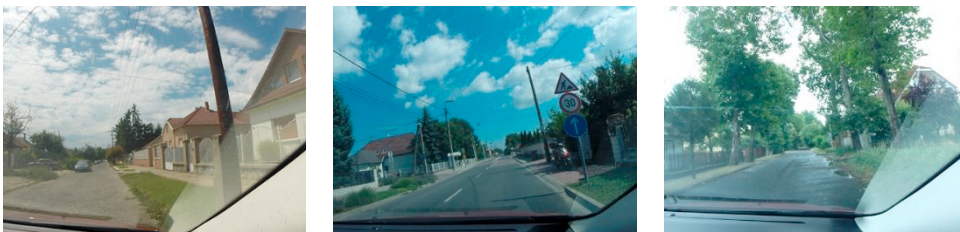Fig. 1. Three downtown areas: one in Vác, one in Százhalombatta and one in Csepel.

Fig. 2. Three residential areas: one in Vác, one in Százhalombatta and one in Csepel.

Fig. 3. Three industrial/commercial areas: one in Vác, one in Százhalombatta and one in Csepel.

Typical images from these urban settlements are shown in Figs. 1 - 3. For each settlement, an example for each of the three urban environments considered in this study is presented. These urban environments are as follows:

- downtown area — featuring one-, or multi-storey buildings built next to, or very close to each other.
- residential area — featuring green spaces and one- and two-storey buildings with somewhat more space between neighbouring buildings.
- industrial/commercial area — featuring factory buildings, workshops, stores with rather spacious yards, as well as supermarkets and bigger shops with parking lots.

A tablet-based Android application – developed earlier for geo-tagged data recording – was adopted for entering the traffic sign and urban environment data manually, while the trajectory data collection was carried out automatically by the application. The data collection personnel involved consisted of two persons: a driver and a data entry assistant.

The driver drove the car along a pre-defined route, while the data entry assistant dealt with the data-entries. To properly document the data collection and to provide means for validating the collected data, a forward viewing camera was installed onto the windscreen to record the road and its surroundings in time-lapse mode. Starting and stopping the photo recording were also the data entry assistant's tasks.

5. Detecting the change of the urban environment from traffic sign data

Based on the collected traffic sign data, the four traffic sign types shown in Fig. 4, namely the 'Maximum speed 30 km/h' signs, the 'Give way' signs, the 'Parking lot' signs and 'No stopping' signs were considered useful for distinguishing the above environments and to detect the transitions between them. The utility of these signs rests on their frequent occurrences in each of these environments, as well as on the sizeable differences between the respective frequencies. Both aspects can be verified by analyzing Fig. 4. The recorded logs of the traffic signs are seen as realizations of a marked inhomogeneous Poisson point process, which herein was approximated – in discrete space – with a marked inhomogeneous binomial point process, as described in Subsection 3.4. A subset of logs were used to tune the parameters of the process models incorporated in the change detectors, while another subset of logs were used for testing. The implementation presented herein breaks up the vehicle trajectory into 50 m long segments and uses six PHCDs – i.e., one for each kind of the transitions – for detecting change along the route. The considered environmental transitions are as follows:

- from residential area to the downtown (Res → Dt), and vice versa (Dt → Res),
- from downtown to industrial/commercial area (Dt → Ind), and vice versa (Ind → Dt), and
- from industrial/commercial area to residential area (Ind → Res), and vice versa (Res → Ind).

The PHCDs are triggered by the lack of – i.e., $(1 - \xi_j)$ in (2) – and the occurrence of a traffic sign – i.e., ξ_j in (2) – respectively, in each segment. (The segments are indexed with j.) The parameters of the PHCDs must be tuned to the respective probabilities in the manner described in Subsection 3.4 and more thoroughly in Fazekas et al. (2016). These probabilities – calculated from the collected data – are shown in Fig. 4.
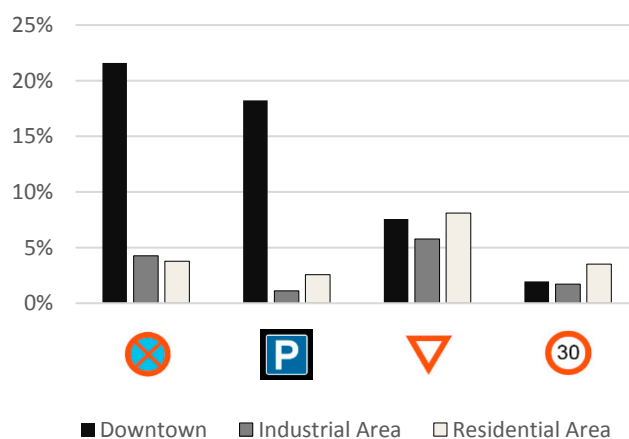

Fig. 4. Probability of the indicated four traffic signs' occurrence – along a 50 m path-length – in downtown, industrial/commercial and residential urban environments based on all the relevant traffic sign data that had been collected from Csepel, Vác and Százhalombatta.

Based on the diagram in Fig. 4, one expects that downtown areas will be fairly easy to distinguish from the other two environments, as the traffic signs located in the former areas have a rather distinctive distribution – over the traffic sign categories – compared to those corresponding to the other two urban environment types.

On the other hand, the industrial/commercial and the residential areas have quite similar distributions over the traffic sign categories. This may indicate that it will be more difficult and will require considerable delay – both in terms of space (i.e., covered driving distance) and time – to distinguish between these two environments.

5.1. Change detection results

Three concrete urban environmental changes – monitored by PHCDs that have been tuned to the probability laws associated with different environmental transitions – are shown in Figs. 5 - 7. The specific urban environmental transitions presented in these figures are (Res → Dt), (Dt → Ind) and (Res → Ind), respectively. The true transitions – as judged and recorded by the data entry assistant – are indicated by the gray shading. The path-lengths are calculated from the start of the given data collection trip. The actual traffic signs locations along the route – for the traffic sign types shown in Fig. 4 – are marked with small red triangles along horizontal axes.

A change is detected when both of the following two conditions are satisfied: the detector output is stalled at some low level (i.e., near zero), and it rapidly – but not necessarily monotonously – increases thereafter. See Subsection 3.4. for details of change detection using a PHCD.

According to the diagram shown in Fig. 5, an urban environmental change, namely a residential area to downtown area transition (i.e., Res → Dt) was detected at the road location corresponding to 60.5 km path-length. This location is situated in the town of Vác. This specific PHCD was tuned to detect such an urban environmental transition based on traffic sign data collected in Csepel (i.e., in another urban settlement). According to the figure, the true transition, however, happened somewhat earlier. The detections shown in Figs. 6 and 7 can be analyzed in a similar manner.
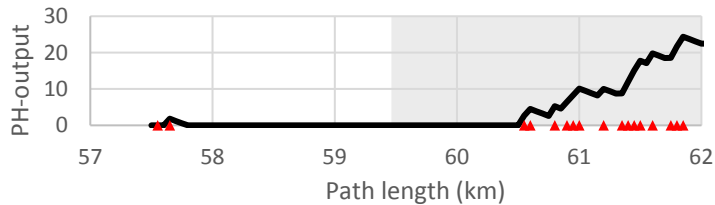

Fig. 5. The output of the PHCD for a Res → Dt transition within Vác
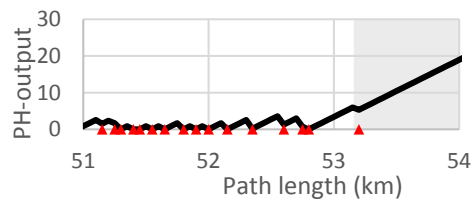

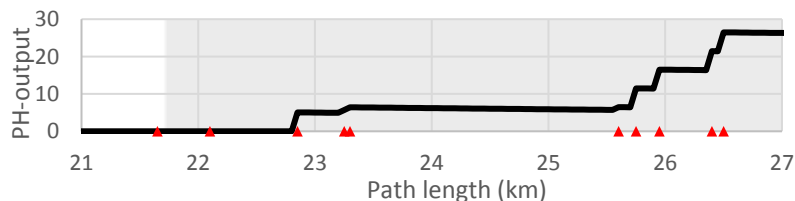Fig. 6. The output of the PHCD for a Dt → Ind transition within Csepel.

Fig. 7. The output of the PHCD for a Res → Ind transition within Csepel.

6. Conclusions

When driving from the periphery of a town to its centre, the urban scenery might change a lot. Such changes can be traced in Figs. 1 - 3. In many cases, the change in the road environment is reflected – and therefore can be traced and detected – in vehicle speeds, in the composition of the vehicular traffic, in air pollution and noise levels, and in the composition and spatial distribution of the traffic signs. Taking logs of traffic signs from a car can be considered

probing the spatial distribution of the traffic signs along the route. Different road environments pose different risks to the car drivers and it is important that the driver is aware of these risks. To assist drivers in this respect, an ADAS function, namely the RoED was proposed and a TSR-based surrogate RoED system was sketched in the paper.

In the presented pilot study, traffic sign logs taken in three Hungarian urban settlements by a human data entry assistant were used to identify (spatial) changes in the urban environment along a route. Three urban environment categories were distinguished both during data collection and in the evaluation phase. The aim was to identify the change points between these environments. The logs are seen as realizations of an inhomogeneous marked Poisson process and the MDL principle was set to work in the form of six PHCDs to identify and locate urban environmental change points from the traffic sign data.

Though the car-based traffic sign data collection covered a much wider range of traffic sign categories, only data for four traffic sign categories were considered relevant for the purpose of environmental change detection and were used to tune the PHCDs. The detectors proved usable in case of each type of environmental transitions considered, though the average delays in detection measured in path-lengths were different. This corresponds to the different Kullback–Leibler divergences of the probability law pairs involved. Generally, transitions involving downtown were detected without much delay as it is shown in Figs. 5 and 6. Note that the diagram in Fig. 6 even shows an early detection, i.e., one that seemingly contradicts causality. This is explained by the somewhat subjective, imperfect and slow categorization of the road environment by the data entry assistant during data entry.

In order to make the TSR-based surrogate RoED technology proposed herein usable in the automotive industry, setting up extensive databases for storing relevant traffic sign and urban environment data would be essential. It is imperative that the data comes from a wide range of regions and countries.

Acknowledgements

References

Amditis, A., Pagle, K., Joshi, S., Bekiaris, E., 2010. Driver–vehicle–environment monitoring for on-board driver support systems: lessons learned from design and implementation. Applied Ergonomics, 41(2), 225–235.

Baikovicius, J., Gerencsér, L., 1990. Change point detection in a stochastic complexity framework. In: the Proc. of 29th IEEE Conf. on Decision and Control, 3554–3555.

Baikovicius, J., Gerencsér, L., 1992. Change-point detection as model selection. Informatica, 3(1), 3–20.

Fazekas, Z., Biró, Zs., Gáspár, P., 2012. Statistical analysis of data describing the relationship between driver, truck and characteristics of the road. Transport Problems, 7(3), 77–85.

Fazekas, Z. Gerencsér, L., Gáspár, P., 2017. Detecting change in road environment via analysis of marked point processes associated with traffic signs. Transportation Research Procedia, 22, 75–84.

Fernandes, L. C., Souza, J. R., Pessin, G., Shinzato, P. Y., Sales, D., Mendes, C., Prado, M., Klaser, R., Magalhães, A. C., Hata, A., Pigatto, D., Branco, K. C., Grassi Jr. V., Osorio, F. S., Wolf, D. F., 2014. CaRINA intelligent robotic car: architectural design and applications. Journal of Systems Architecture 60, 372–392.

Fua, P., Hanson, A. J., 1991. An optimization framework for feature extraction. Machine Vision and Applications, 4(2), 59–87.

Hinkley, D., 1971. Inference about the change-point from cumulative sum tests. Biometrika 58, pp. 509–523.

Kiernan, J. Terzi, E., 2009. Constructing comprehensive summaries of large event sequences. ACM Transactions on Knowledge Discovery from Data, 3(4), 21.

Lakshmanan, L. V. S., Ng, R. T., Wang, C. X., Zhou, X., Johnson, T. J., 2002. The generalized MDL approach for summarization. 28th Int. Conf. on Very Large Data Bases, 766–777.

Lorden, G., 1971. Procedures for reacting to a change in distribution. The Annals of Mathematical Statistics 42, 1897–1908.

Møgelmose, A., Trivedi, M. M., Moeslund, T. B., 2012. Vision-based traffic sign detection and analysis for intelligent driver assistance systems: perspectives and survey. IEEE Transactions on Intelligent Transportation Systems, 13(4), 1484–1497.

Page, E. S., 1954. Continuous inspection schemes. Biometrika 41, pp. 100–115.

Piao, J., McDonald, M., 2008. Advanced driver assistance systems from autonomous to cooperative approach. Transport Reviews, 28(5), 659–684.

Gao, Q., Li, M., Vitányi P., 2000. Applying MDL to learn best model granularity. Artificial Intelligence, 121, 1–29.

Rissanen, J., 1978. Modeling by shortest data description. Automatica 14, 465–658.

Rissanen, J., 1998. Stochastic Complexity in Statistical Inquiry. World Scientific Series in Computer Science, 15.